UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF NEW YORK

CHEVRON CORP.,

Plaintiffs,

-against-

Case No. 1:12-mc-65 GLS/CFH

Hon. Gary L. Sharpe

STEVEN DONZIGER, *et al.*,

Defendants.

DECLARATION OF SETH SCHOEN IN SUPPORT OF MOTION OF NON-PARTY JOHN DOE MOVANTS TO QUASH SUBPOENAS TO MICROSOFT, INC. SEEKING IDENTITY AND EMAIL USAGE INFORMATION

I, Seth Schoen, declare as follows:

1. I am a Senior Staff Technologist with the Electronic Frontier Foundation in San Francisco, California, and I make this declaration on my own personal knowledge. I have worked with computers professionally for over a decade and have testified about communications systems in three courts and before the United States Sentencing Commission.
2. The purpose of this declaration is to provide a general introduction to IP addresses, the use of IP addresses to track location, and how that information could be used to associate a person with others.

Introduction to Internet Protocol Addresses

3. An Internet Protocol address (or "IP address") is a numeric value used to identify the network location of a computer or set of computers on the Internet. Every computer on the Internet needs to have an IP address in order to communicate with other computers on the Internet. Internet routers use the IP address to decide where to send communications

a particular computer user.[1] The address is normally written as four numbers from 0 to 255 separated by dots.[2] For example, one of the web servers operated by the Electronic Frontier Foundation uses the address 64.147.188.11, while the District Court for the Northern District of California's web server uses 206.18.146.127, while the District Court for the Northern District of New York's web server uses 199.107.21.60.

4. IP addresses are allocated to Internet service providers (ISPs) in blocks of consecutive addresses out of a worldwide pool of around four billion possible addresses through geographically based non-profit organizations known as regional Internet registries.[3] ISPs can further delegate these addresses to smaller entities such as businesses, Internet cafés, or smaller ISPs.[4] ISPs can also assign an IP address directly to an individual computer. This assignment process is frequently automated and the assignment can be short- or relatively long-term.[5]

5. Because IP addresses are allocated in this way, they can convey approximate information about a computer's location, how the computer is connected to the Internet, and what individual or entity is using that computer to connect.

6. Individual users connect using different IP addresses depending on where they are. Multiple users who are using the same local-area network can share a single IP address (most often when they use a shared router or wireless connection) and hence appear to connect to the Internet through the same IP address, whether at different times or at the same time.[6] If users of the same local-area network are *not* sharing a single IP address,

---

1 Eric A. Hall, *Internet Core Protocols: The Definitive Guide*, 37-40 (O'Reilly and Associates,

2 *See* Radia Perlman, *Interconnections Second Edition,* 199 (Addison Wesley Longman, 2000). This declaration uses "Internet Protocol address" to refer to addresses using version 4 of the Internet Protocol ("IPv4"), which has been extensively used worldwide since 1980. Due to exhaustion of the pool of distinct IPv4 addresses, the Internet is now in the course of switching to version 6, which uses significantly longer addresses.

3 *See* American Registry for Internet Numbers, "Internet Number Resource Distribution," https://www.arin.net/knowledge/distribution.pdf (last visited Oct. 22, 2012).

4 Hall, *supra* note 1, at 40-41.

5 *See* IP address, http://en.wikipedia.org/w/index.php?title=IP_address&oldid=518867856 (last visited Oct. 22, 2012).

6 *See* Yinglian Xie *et al.*, "How Dynamic Are IP Addresses?", in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer*

and instead have distinct addresses, their IP addresses will generally be numerically adjacent (with the beginning portion of the address identical, and only the final portion different).[7]

7. An IP address may identify the network through which a network-enabled device (such as a desktop computer, laptop computer, tablet or smartphone) is accessing the Internet. When a portable device moves from an Internet connection on one network (the network connection at one's home, for example) to an Internet connection on another network (a local coffee shop), the IP address associated with the portable device changes to reflect that the device is connected to that particular network. This change is normally carried out completely automatically and is transparent to the user.[8]

8. Many host computers of websites, including the operators of popular web-based e-mail services like Yahoo! Mail, Gmail, and Microsoft Hotmail, maintain logs that list the IP address of visitors along with date and time information. Websites that utilize a log-in feature typically maintain a log of IP addresses and other data associated with the particular user who logged in, such as the date and time of log-in and the duration of time the user visited the website. If a user accesses the website with a portable device from different locations, then the log data about that user will include a variety of different IP addresses. Because of the way they were assigned, these different IP addresses will reflect the location and movement of the device and its owner.

Communications, https://research.microsoft.com/pubs/63680/sigcomm07-onefile.pdf (last visited Oct. 22, 2012); Jeff Tyson, "How Network Address Translation Works," http://computer.howstuffworks.com/nat.htm/printable (last visited Oct. 22, 2012).

[7] *See* "Subnetwork," http://en.wikipedia.org/w/index.php?title=Subnetwork&oldid=517971549 (last visited Oct. 22, 2012); "Classless Inter-Domain Routing," http://en.wikipedia.org/w/index.php?title=Classless_Inter-Domain_Routing&oldid=518823902 (last visited Oct. 22, 2012).

[8] Today, the automatic assignment of a new address would be handled by the Dynamic Host Configuration Protocol (DHCP). "Dynamic Host Configuration Protocol," http://en.wikipedia.org/w/index.php?title=Dynamic_Host_Configuration_Protocol&oldid=519194559 (last visited Oct. 22, 2012).

Using IP Addresses and Associated Information to Determine Location

9. A large amount of data accumulated over a lengthy period of time that includes IP addresses and dates and times of usage sessions—as one might get from a heavily trafficked and frequently used web service such as an email provider—can readily present a detailed picture of a person's movements from one location to another, especially if that person is an avid laptop or tablet user.

10. For instance, a laptop will receive a different IP address when it connects to the Internet from different locations.[9] If a laptop's owner uses the machine from her workplace in the morning, a café in the afternoon, and her home in the evening, she will present at least three different IP addresses over the course of a single day. A traveler who brings a laptop to a different country and goes online there will receive an IP address unrelated to the IP address he used at home.

11. The WHOIS service, which can be accessed through web sites such as http://www.domaintools.com, is a public database that permits a user to find out to whom a regional Internet registry has allocated a block of IP addresses. A user can input a numerical IP address and obtain the registry's information about the assignee of that IP address, which might be an ISP or other entity. When IP addresses have been allocated directly to an organization that makes use of them, WHOIS can sometimes associate an IP address with an exact physical location. For example, inputting the IP address 156.128.118.200 into WHOIS shows that the number is associated with the Administrative Office of the U.S. Courts, which is located at One Columbus Circle NE, Washington D.C. On most occasions, however, WHOIS will only associate an IP address with an Internet service provider's office, which is often in the same region as its subscribers or users but does not reveal their exact locations.

12. WHOIS records and other information sources provide geographic information about where ISPs operate and where they use particular ranges of IP addresses. This means that servers located in California typically have IP addresses traceable to California, servers

---

[9] *See* University of Illinois Campus Information Technologies and Educational Services, "Network Access While Traveling," http://www.cites.illinois.edu/network/access/travel.html (last accessed Oct. 22, 2012).

located in New York typically have IP addresses traceable to New York, and servers in Ecuador typically have IP addresses traceable to Ecuador. This geographic location information is generally publicly available. For instance, the IP address 199.83.220.233 is easily traceable to San Francisco through free websites available to the public such as www.geobytes.com. Even when location information is not publicly available, a subpoena to an ISP can generally elicit the specific geographic location for a particular IP address.

13. Where that is the case, IP address records can still be used in the service of pinpointing a person's location and movements, particularly in the context of litigation where parties can exercise subpoena power. Internet service providers typically maintain records of the physical addresses associated with a given subscriber, for both billing and service purposes, and typically also record historical information about which subscriber used a particular IP address. Once the a litigant has a list of IP addresses, it can subpoena subscriber information from the corresponding ISPs for the specific IP addresses and develop a detailed picture of a person's location and movements from that subscriber information.

14. An IP address may reflect the place where a person accesses a certain Internet service. This information might demonstrate that a person accessed the Internet from a certain physical location, like a building or even particular organization's office. (The IP address ranges used by particular organizations are not necessarily published, but it could sometimes be possible to determine or recognize them on the strength of other records that incidentally reveal them, or via subpoena.)

Using IP Addresses and Associated Information to Determine Associations With Other People

15. A large amount of data accumulated over a lengthy period of time that includes IP addresses and dates and times of usage sessions—as one might get from a heavily trafficked web service such as an email provider—can reveal a person's physical proximity to other Internet users who may share the same IP address. This information could be used to map a person's associates.

16. If Internet usage records showed that two individuals were accessing the Internet from

the same IP address (or numerically proximate IP addresses) on a particular day and time, this would tend to show that they were accessing the same Internet network at the same time. This would strongly suggest they were in the same physical location at the same time and could create a reasonable inference that they met with one another.

Consequences of Disclosure of Long-Term IP Address Records

17. Chevron seeks a nine-year span of IP logs associated with 101 email accounts. If it were made available to Chevron, this information would tell Chevron when the targeted individuals were in the United States or abroad. It would tend to show when they where in a particular town, and when they were at home or at work. It could be used to determine when they visited an office of a particular organization, and when each email account holder was in the same place at the same time as other individuals whose IP addresses have been revealed, potentially meeting with each other.

18. The information Chevron seeks can also reveal intensely personal details about the account holders' lives. A habitual e-mail user might check a given e-mail service multiple times per day, and a long-term view of data about this use could evince quite significant facts relating to work habits, personal relationships, and changes in someone's employment or living situation. For example, if the IP logs show that a person signed into his email account from an IP address associated with another person's home, that information suggests he visited that home. If he signs into his email account using that same IP address late at night and again the following morning, it creates a reasonable inference that he spent that night at that home and may have an intimate relationship with that person who lives there. If he repeats this pattern over time, it might suggest that relationship is a serious one.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct to the best of my knowledge. Executed on October 22, 2012.

Seth Schoen